



Extracting Key Content



from Images

A new system called the Image Content Engine is helping analysts find significant but often obscure details in overhead images.

ADVANCEMENTS in imaging technologies, particularly in remote sensing systems, are producing vast amounts of data that can easily overwhelm human analysts. A team of Livermore engineers, computer scientists, and physicists has come to the aid of overburdened analysts who need to quickly analyze large volumes of overhead images. The team's new extraction system, called the Image Content Engine (ICE), allows analysts to search massive volumes of data in a timely manner by guiding them to areas in the images that likely contain the objects for which they are searching.

ICE was developed under a three-year Laboratory Directed Research and Development (LDRD) Strategic Initiative begun in 2003. It encompasses a new approach for the computer-aided extraction of specific content information from different kinds of images, but especially those images taken with overhead sensors.

Imagery analysts in the Laboratory's Nonproliferation, Homeland and International Security (NHI) Directorate are using ICE to identify objects such as specific types of buildings or vehicles. A


variation of ICE is being evaluated for adoption by the National Ignition Facility's (NIF's) Optics Inspection Analysis Team. In addition, the system may be applicable to other fields, including nondestructive analysis, biological imaging, astronomy, and supercomputer simulations of nuclear weapon performance.

"We need new ways to handle the increasingly vast amounts of data generated by sensors," says engineer David Paglieroni, technical lead and co-principal investigator of ICE. Paglieroni, who leads the Engineering Directorate's Imagery Sciences Group, notes that information overload is a serious problem in the intelligence community because analysts face a flood of data from different sources.

Jim Brase is the other ICE co-principal investigator and head of the Optical Science and Technology Division of the Physics and Advanced Technologies Directorate, a group that develops advanced detectors for astronomical research and national security. The ICE team also includes engineers Barry Chen, Chuck Grant, Aseneth Lopez, Doug Poland, and George Weinert; computer

scientists Jim Garlick and Marcus Miller; and postdoctoral researchers Siddharth Manay and Faranak Nekoogar.

Brase notes that the laboratory has been a leader in data mining, which involves finding items of interest in large amounts of data. For example, Livermore computer scientists are helping to tease out the most relevant features in three-dimensional visualizations of scientific simulations. (See *S&TR*, November 2004, pp. 12–19.) Livermore researchers have also developed several algorithms to extract items of interest in supercomputer visualizations of simulated weapons performance. The goal of these efforts is to enable analysts in all of the Laboratory's programs to focus on the most important details, which are often buried within massive amounts of visual information.



Extracting Desired Content

The ICE development team worked closely with imagery analysts in NHI's Z Division (International Assessments and Knowledge Discovery). August Droege, leader of the Precision Intelligence Group, explains that analysts pore over large amounts of imagery, either at light tables or at computer workstations, often looking for obscure objects. "Our analysts are very fast, but there can be thousands of images to sort through," says Droege.

The resolution and the amount of area covered in an image can vary widely, with many images covering enormous areas. For example, a commercial satellite image with 1-meter-per-pixel resolution (6,000 rows by 10,000 columns of pixels) covers an area of about 60 square kilometers. Wes Spain, leader of Z Division, says, "Efforts over the past 20 years aimed at analyzing this kind of data with computers have had limited success."

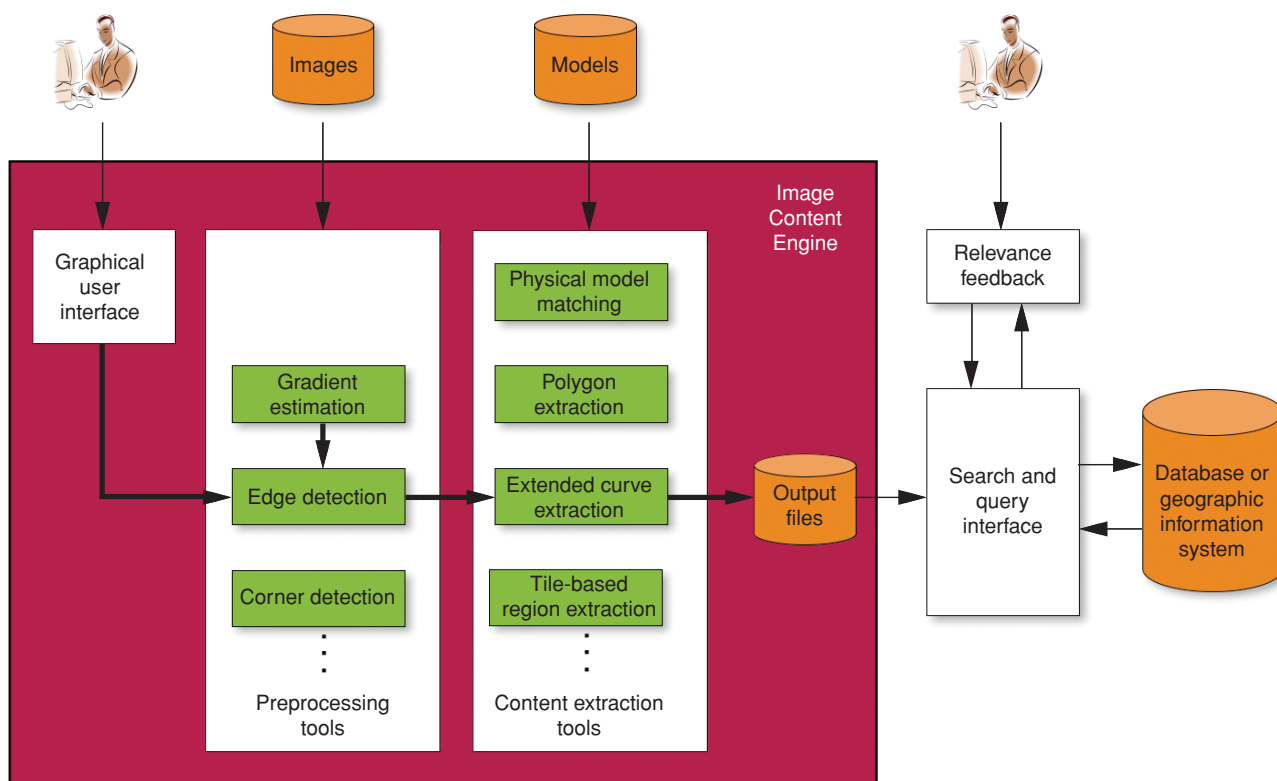
ICE can accommodate images acquired with different types of overhead sensors and at varying resolutions. The software is also able to account for the fact that images are taken at different times of the day, during different seasons, and under changing weather conditions. Images often contain distracting background clutter, and potential objects of interest can be in full or partial shadow, occluded by trees, or obscured by snow or clouds. "Finding specific objects, such as particular types of buildings or vehicles, in overhead images that cover hundreds of square kilometers, is a difficult task," says Paglieroni.

The ICE architecture can run on different computing platforms and operating systems, such as Windows or Linux operating systems, laptops or powerful clusters of computers, and isolated or networked processors. The ICE software contains a library of algorithms, each of which focuses on a specific task. The algorithms can be

chained together in pipelines configured through a graphical user interface. Each pipeline is designed to perform a specific set of tasks for extracting specific image content.

One of the most useful ICE algorithms, known as gradient direction matching (GDM), was developed by Paglieroni. GDM uses a novel approach for implementing equations to rapidly "pull" objects of interest out of images that have large amounts of visual clutter. The algorithm compares pixel gradient directions (the direction of flow from dark to light) in an image being analyzed to pixel gradient directions that are perpendicular to unoccluded edges in submitted models. As a result, GDM is relatively insensitive to image brightness and contrast variations. For overhead images, ICE uses GDM for matching objects to a variety of models, for extracting vertices (where two or more lines meet) of specified polygons, and for computing certainties associated with

The Image Content Engine (ICE) software contains a library of algorithms, each of which focuses on a specific task. The algorithms can be chained together into pipelines configured through a graphical user interface. In this example, a user requests ICE to extract roads (extended curves) from images.




matches to submitted polygons. GDM has also been successfully used in NIF optics inspection to search for diffraction patterns that signal defects. “GDM has been a gold mine for us,” says Paglieroni. “We consider ourselves very fortunate when one algorithm can be used to support more than one important processing task.”

ICE also provides tools for extracting regions, extended curves, and polygons from images. The region extraction tool breaks images into small, adjacent square tiles containing one or more pixels. For each tile, the algorithm searches for spectral or textural characteristics and then groups tiles with similar features into regions. This tool is useful for separating distinct areas such as forests, bodies of water, plowed fields, and clusters of buildings from the image background.


The extended curve extraction tool is used to find lines of communication, such as roads, power lines, and canals, in overhead images. ICE uses a hierarchical approach in which mistakes made in processing at the pixel level are corrected at successively higher levels, such as the line-segment or curve levels. The ICE team developed a novel approach to consolidate collections of broken line segments that are nearly parallel into single, consolidated line segments. These segments are then assembled into consolidated curves. This approach, which is still under development, allows ICE to extract lines of communication from even highly cluttered scenes, a capability that other image-analysis tools rarely possess. (See the figure on p. 16.)

The polygon extraction tool, which is also under development, is useful for automatically extracting man-made objects, such as buildings and moving vehicles, from overhead images. The ability to perform such automated extractions facilitates computer-aided searches for objects in overhead images. The tool assigns each polygon vertex a pixel location, sharpness, and orientation. By analyzing groups of vertices, the algorithm


(a)



(b)



(c)



The gradient direction matching (GDM) algorithm compares (a) the edges of a projected model to (b) a small section, or block, of a much larger overhead image. (c) The small colored dots indicate the best local matches, which are then sorted and ranked in order for inspection by an analyst.



ICE extracts a region composed of tiles (green boxes) whose features are consistent with built-up areas.

quickly extracts polygons of prescribed geometry, independent of their position and orientation in the image. If the model specifications are relaxed, polygons of arbitrary size or with a particular ratio of width to height can be found. “We are encouraged that GDM appears to be outstanding at extracting vertices, which

has traditionally been a difficult task,” says Paglieroni. (See the top figure on p. 17.)

The ICE team envisions that by combining diverse content extraction tools, analysts will be able to search for complicated patterns containing different pieces of content. As an analyst adds pieces of content to be extracted from an

image, the false-alarm rate will drop. That is, fewer areas in the image will be mistakenly identified as a feature of interest.

Specifying and Matching Models


ICE provides an interface for specifying the sets of images to be processed and how they are to be processed. A user either specifies the model of an object to search for from a registry furnished by ICE or supplies the model to match. The user-supplied model can be a specified polygon or a physical model with two or three dimensions (in meters) based on a detailed drawing or, alternatively, with dimensions derived from a reference photograph.

ICE divides images into overlapping blocks, and the amount of overlap depends on the size of the object being searched. At each pixel location, the edges of the submitted model are projected onto blocks at typically 75 different orientations. The gradient direction perpendicular to the edge boundary of the projected model is computed for each pixel in each block using GDM. The best match over all orientations at every pixel is saved for subsequent ranking. “This processing stage is computationally expensive and may require parallel-processing compute clusters,” says Weinert. With adequate computational horsepower, ICE can process hundreds of images in a few hours.

In the search and query stage, ICE creates image thumbnails (small image blocks) of the closest matches to submitted models. These matches are sorted in order of decreasing similarity to the model and presented to an analyst for visual inspection and interpretation. Much like an Internet search engine, the program assigns each match a score, ranging from 0 to 1.00 (0 to 100 percent). Although parallel computing may be required for ICE to process large sets of images in a timely manner, ranked thumbnails can be generated quickly on a personal computer.



When an analyst clicks on a thumbnail, the surrounding context for that thumbnail

An example of extracting extended curves from an image shows (a) traditional edges, as might be found with other image-processing systems and (b) consolidated lines and curves.

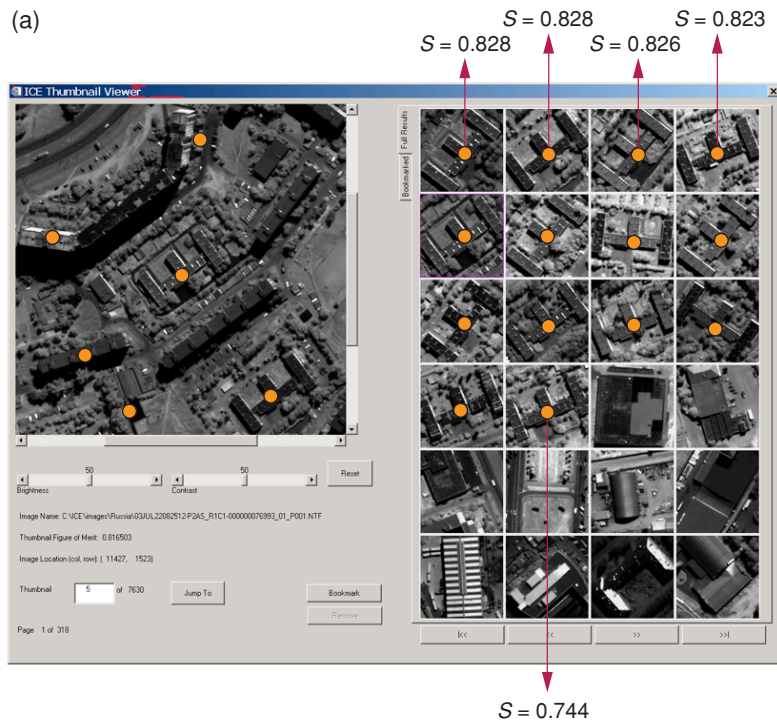


is displayed in a larger window. “The object of interest may not be contained in the thumbnail with the highest score, but the probability is high that the object of interest will be present among relatively few top-scoring thumbnails,” says Weinert. Because analysts can miss targets in thumbnails that are too dark or light or that have poor contrast, a user can enhance the brightness and contrast settings. For example, a user can brighten an object in shadow or darken an object in snow.

ICE can compute geographic coordinates (latitude and longitude) for matches extracted from images. The ICE team is developing the capability to link those matches to geographic information system data such as street and topographic maps. Also in the works is the ability to store ICE matches in commercial databases and semantic graphs (advanced techniques that capture complex relationships among detected objects of interest) along with nonimage data generated by other systems.



(a) An example of 90-degree polygon corner extraction using the GDM algorithm shows how the algorithm detects the flow of light to dark gradient magnitudes in the image and (b) matches the gradients to a model of 90-degree polygon corners. (c) Polygons in the image are then extracted from the candidate corners.



(a) In the search and query stage, ICE creates image thumbnails of the closest matches to (b) a submitted model. These matches are sorted in order of decreasing similarity to the model and presented to an analyst for visual inspection and interpretation. In this example, the upper left-hand thumbnail has a similarity (S) of 0.828 (the highest of any match). An analyst can click on any thumbnail and see it in a larger context. The thumbnails can also be adjusted for brightness and contrast. (c) Each match is captured as a vector and stored in an output file of matches.

ICE Proves Itself

For more than a year, ICE has demonstrated significant potential to increase productivity by focusing imagery analysts' attention on possible objects of interest. "ICE gets the eyes of analysts on the highest priority images," says Spain. "It serves almost like a triage function, telling us which images are the most deserving of our attention."

The program can also be a huge timesaver. "We can potentially do so much more with ICE," says Droege. She estimates that ICE could reduce more than one week of intensive analytical work into a couple of hours. "ICE significantly increases the probability of finding something of interest. It does the heavy lifting for us," she says.

Droege is especially appreciative of the program's flexibility. Analysts can ask the program to look for a specific structure that matches exact measurements or for one that merely suggests an object such as an airfield. "We may want to search just for a structure in a desolate area where one might not expect anything to be built," she says.

Paglieroni explains that ICE provides a computer-assisted approach to analysis. In contrast, the computer-automated approach attempts to replace human abilities for analyzing and interpreting images with computer programs. However, computers are not good at interpreting images. "The computer-assisted approach uses the strengths of computers to enhance the strengths of human analysts," he says.

Scrutinizing NIF Optics


The NIF Optics Inspection Analysis Team in conjunction with the ICE collaborators has adapted the GDM algorithm to strengthen NIF's computerized optics inspection system. The project, called Finding Rings of Damage in Optics (FRODO), is an Engineering Directorate technology-base project. Technology-base projects adapt new techniques to the specific needs of Livermore programs.

NIF currently has two bundles of eight laser beams each, with a third bundle undergoing installation. When completed, NIF will have thousands of square-shaped optics that guide, amplify, and focus light from 192 laser beamlines (24 bundles)

onto a tiny target. Camera systems throughout the laser take pictures of individual optics as part of the NIF Optics Inspection Analysis System, which analyzes the images to automatically monitor the condition of optics. Up to 80 images are taken of the optics in each laser beamline every time a full inspection is requested. When all 192 beams are operational, some images will be required multiple times per day. The goal is to automatically detect and characterize changes over time, providing information to help managers decide whether to repair, replace, or continue using an optic.

Engineer Laura Kegelmeyer, assigned to NIF from the Engineering Directorate's Signal, Image Processing, and Control Group, is the principal investigator of the FRODO project. Kegelmeyer says that NIF optics must be frequently inspected to operate safely, ensure quality performance, and determine when optics need to be refurbished. "The slightest imperfection can affect the uniformity of the laser beam, and some imperfections can grow under repeated exposure to laser pulses," says Kegelmeyer. "We want to find any defect less than 500 micrometers in diameter so

The GDM algorithm has performed extremely well in detecting diffraction patterns in test images of small flaws in laser optics. (a) In the test image, a faint diffraction pattern is detected. (b–c) To detect the diffraction rings, the operator uses a model that produces a field of light-to-dark gradient vectors flowing away from the center of a circle. (d–e) The GDM algorithm then identifies regions in the test image whose light gradient directions best match those of the luminance disk. After GDM identifies potential rings, ICE ranks the rings by how well they fit the model.



the optic can be refurbished before the defect grows too large.”

Custom algorithms developed at Livermore currently detect and characterize defects on NIF optics using a “direct” approach to find irregularities in images of the optics. However, the location of some optics limits the detail that can be obtained in direct images.

FRODO is a complementary approach that searches for indirect evidence of defects in the form of diffraction ring patterns. These circular patterns of light appear on images of optics that are located downstream from the optics with imperfections. “The situation is analogous to dropping a pebble into a pond,” says Paglieroni. “The larger the pebble, the stronger the ripples.” The pebbles are defects that occur in upstream NIF optics but which manifest themselves as diffraction patterns closer to the camera. The further away from the defect, the larger are the concentric ring patterns.

The GDM algorithm has performed extremely well in analyzing images of optics and detecting diffraction patterns caused by distant flaws. NIF engineers are hopeful that this indirect approach will complement their direct inspection approach by helping detect flaws that are outside the camera focal range or for which image resolution is limited.

When attempting to detect diffraction rings, the FRODO software invokes the GDM algorithm by using a luminance disk, which is a model that produces a field of light-to-dark gradient vectors all flowing toward the center of a circle. This model is similar to those used by ICE when searching for specific objects in overhead images but has a round shape rather than a shape similar to a road or building.

The algorithm’s final output is a set of estimates for ring locations on the image as well as the estimated size of the upstream defect and the distance to the optic. Kegelmeyer says, “Our goal is to demonstrate the robustness of the algorithm in real and simulated images and in images that have no defects, single defects, or multiple defects of different sizes. We have found that the algorithm can consistently and efficiently identify the associated diffraction patterns.”

In certain instances, however, the algorithm encounters problems. Extremely small diffraction rings are difficult to detect as are rings cut off near the edge of an image when two rings overlap more than 50 percent. The FRODO team has optimized the algorithm parameters and the postprocessing steps to minimize the number of false alarms. “These false-alarm mitigation techniques look promising,” says Kegelmeyer. Once this issue is resolved, the GDM algorithm can be incorporated into the suite of NIF optics inspection analysis codes.

Building on Success

Livermore imagery analysts’ success with ICE has sparked interest by several federal agencies involved in national security. In the meantime, imagery analysts are pushing ICE to its limits by searching for increasingly smaller objects.

ICE has potential applicability to other areas of experimental science, including physics, biology, and environmental science, in which mining massive archives of complex measurement data is an essential research activity. One possibility under consideration involves using ICE for automatic extraction of data generated by

the Large Synoptic Survey Telescope (LSST). (See *S&TR*, November 2005, pp. 12–18.) The telescope, scheduled for completion in 2012, will provide digital imaging of objects in deep space across the entire sky. LSST will create 24 gigabytes of data every 30 seconds, unprecedented in astronomical data gathering. Effectively managing the vast amount of LSST data is the most challenging aspect of the project. Computer scientists Miller and Garlick are involved in developing computational approaches for automatically mining the thousands of digital images that LSST will record daily.

A new LDRD Strategic Initiative, called Predictive Knowledge Systems (PKS), is building, in part, on the capabilities of ICE. PKS will pull together multiple sources of information, such as imagery, radio intercepts, and other sensor data, and correlate the data in space and time. Brase, who is principal investigator for PKS, says the project is aimed at nuclear nonproliferation and homeland security applications. In disciplines where the sheer amount of information can be overwhelming, ICE is providing a way to find the proverbial needle in a haystack. In this case, finding that needle may save lives.

— *Arnie Heller*

Key Words: gradient direction matching (GDM) algorithm, Image Content Engine (ICE), Large Synoptic Survey Telescope (LSST), National Ignition Facility (NIF), Predictive Knowledge Systems (PKS).

For further information contact
David Paglieroni (925) 423-9295
(paglieroni1@llnl.gov).